A Learned Approach for Lump Identification In Soft Tissue via Palpation

Kirk A. Nichols
Department of Mechanical Engineering, Stanford University
Kirk.nichols@stanford.edu

ABSTRACT

With the increase in robot-assisted minimally invasive surgeries, restoring haptic information lost by the lack of the surgeon's direct contact with the patient is crucial. Here I seek to restore lost stiffness information by gathering force and position data taken from a robot palpating an artificial tissue embedded with a stiff object simulating a circular tumor. Using machine learning algorithms, I estimate the location and radius of the tumor. Results show that of the learning methods explored, the supervised learning methods perform better than their unsupervised counterparts, with an average error of false positives plus false negatives of 10%. However, the most robust method of these uses a supervised learning algorithm to first threshold the palpation data based on stiffness and then estimating the radius of the tumor using a second machine learning algorithm. I intend to continue and expand this work in the CHARM lab, PI Allison Okamura.

1 INTRODUCTION

In the medical community, tactile sensations provide an array of diagnostic capabilities. Cancerous regions in soft tissues like in the prostate and breast tissues, for example, typically manifest themselves as hard lumps and are significantly harder than the surrounding tissue. During open surgery, surgeons can directly contact and palpate the tissue. Minimally invasive surgery (MIS) and robot-assisted minimally invasive surgery (RMIS) are becoming increasingly popular as they provide increased dexterity and control for surgeons while reducing patient trauma [1]. However, the lack of tactile feedback limits surgeons' capabilities, and it has been shown that RMIS presents with a greater likelihood of leaving behind cancerous cells upon biopsy of a diseased region [4]. Finding some method to convey the shape and location of a suspected lump would restore some of the information lost with the use of minimally invasive surgical methods.

I investigated use of a robotic system with an attached force-torque sensor to provide a data set from which machine learning tools can identify the size and location of a lump in an artificial tissue sample. Section 2 describes the experimental setup used, including the hardware and software methods used to gather data. Section 3 describes the experimental methods used, and introduces five learning methods tasked with identifying the shape and size of the lump in the artificial tissue. Results are presented in Section 4 and discussed in Section 5. Future work and conclusions drawn from the experiment comprise of Sections 6 and 7, respectively.

2 EXPERIMENTAL SETUP

The setup of the experiment involved creating a platform where upon palpating an artificial tissue, information regarding the stiffness of the palpation location as well as the location itself can be gathered.

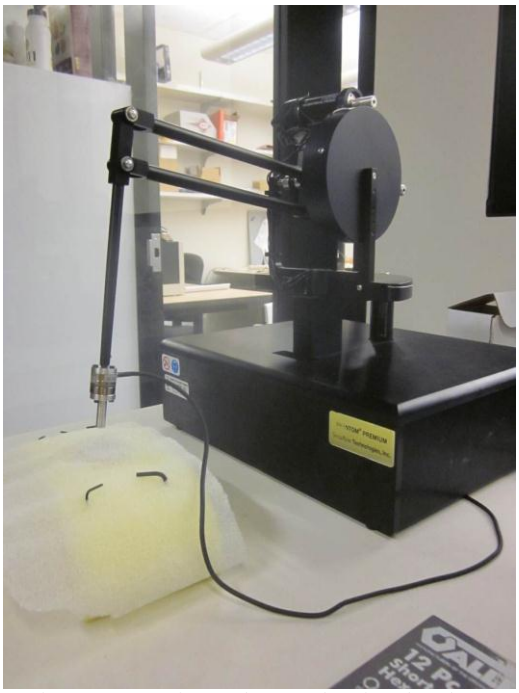

Figure 1: Experimental setup showing the Phantom Premium, the Nano 17 force-torque sensor, and the artificial tissue.

2.1 Hardware

The hardware used for this project are a Phantom Premium 1.5A (Sensable Technologies, Inc.) haptic device and a Nano 17 force-torque sensor (ATI Industrial Automation). The Phantom Premium in its current configuration is a 3-DOF desktop haptic device (Figure 1). The Nano 17 force-torque sensor measures forces and torques along all three axes. However, for the purpose of this project only the force along the z-axis of the sensor (pointing into the end effector of the Phantom Premium as pictured in Figure 1) was gathered.

2.2 Artificial Tissue

Figure 2: Close-up of force-torque sensor with uncovered artificial tissue

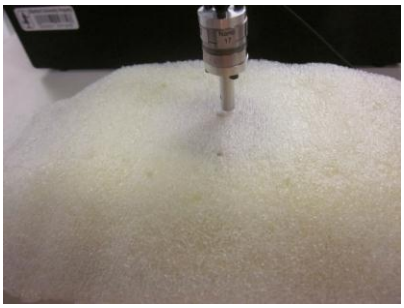

Figure 3: Force-torque sensor with covered artificial tissue

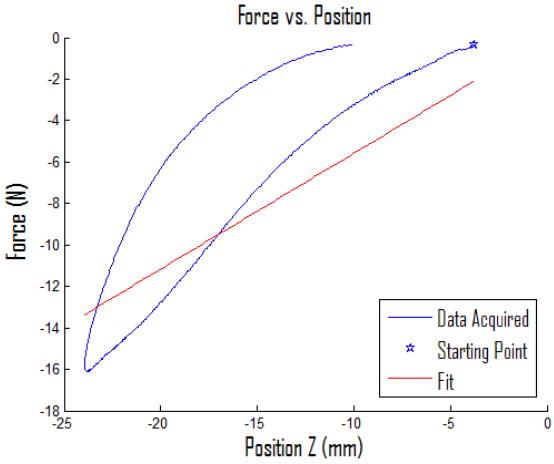

Figure 4: A typical Force vs Position curve collected from a single palpation.

The artificial tissue sample consisted of a sponge with a hard, circular object embedded inside (Figure 2). Circular objects are commonly used to approximate tumors [2]. This sample was covered with a thin piece of packaging material to add more variance between the stiffnesses of the embedded object and surrounding sponge (Figure 3).

2.3 Data Collection

The data sets were gathered by manually palpating the artificial tissue with the robot arm. The robot's encoders were used to gather data about the position of the end effector. Figures 5 and 6 show the different sets of data varied with respect to the rotation of the artificial tissue.

2.4 Tumor Registration

For a ground-truth measurement of the location and size of the tumor, the tumor shape was traced by the end effector of the robot and subsequently registered. Points within the registered shape would be classified as part of tumor and points outside the registered region would be classified as not part of the tumor.

3 Experimental Methods

Five different methods were employed to predict the shape of the tumor. The first three of these methods fall under supervised learning, while the last two are unsupervised, although the fifth method does utilize the notion of a "training set", the labeling for this set is constructed in an entirely unsupervised fashion.

For all the methods, an initialization step was first performed to convert the force/position data for each individual palpation into a stiffness value that would associate with an XY location on the sponge.

3.1 Initialization

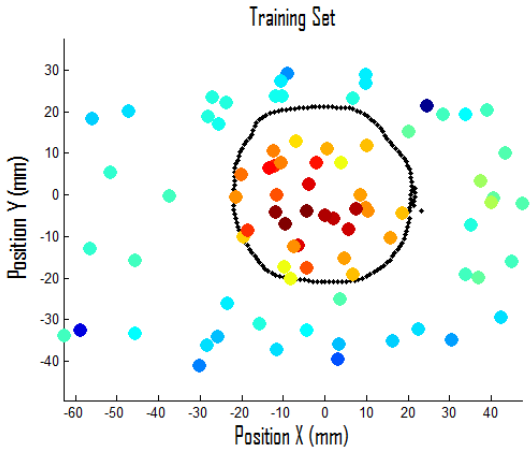

Figure 5: Training Set Used for All Algorithms. Stiffer points are shown in red, while softer points are in blue. The black circle is the trace of the circumference of the tumor taken during registration.

For each palpation, the force-position data acquired from the Nano17 was fitted by least squares to a straight line, with the assumption that the palpated environment can be well approximated by a linear spring. This model was chosen for its simplicity and it appears to capture the features of the resulting Force vs Position plot well (Figure 4). The slope of the line is the stiffness value used for that individual palpation.

3.2 Learning Algorithms

Driving the learning behind the various learning methods described in this section are three different learning algorithms, Logistic Regression (LR), Gaussian Discriminate Analysis (GDA), and a Support Vector Machine (SVM). The convex optimizer used to solve the SVM with regularization equations is a MATLAB toolkit developed by Michael Grant and Stephen Boyd called CVX [3]. The SVM regularization weighting parameter 'C' was tuned by hand.

3.3 Supervised Learning Training Set

The training set used for all the learning algorithms is pictured in Figure 5. The data in this set was found to have the least training set error among all the data sets acquired. Note that in Figure 5, a data point is stiffer if its color is more towards the red color spectrum of the scale, while softer points are depicted towards the blue color spectrum. The three test sets are seen in Figure 6.

3.4 Stiffness Learning followed by Radius of Tumor Learning

The first method used the training set to develop a threshold for labeling a point as in the tumorous region or not in the tumorous region based on the observed stiffness value. The resulting threshold would take a data point in the test set and conclude that

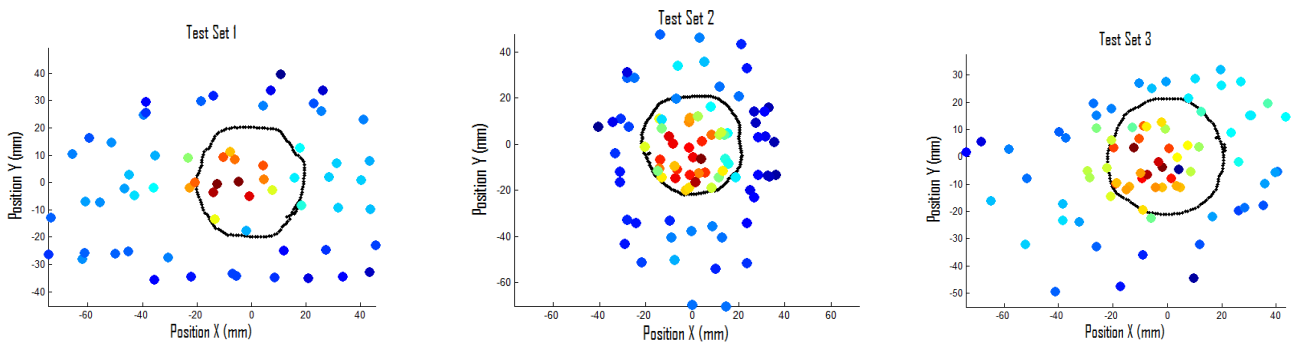

(a) (b) (c)

Figure 6: Test sets used. Note, the black circle, depicting the shape of the tumor, is what we seek to estimate, however they're drawn here for convenience. Note that none of these sets share the same orientation.

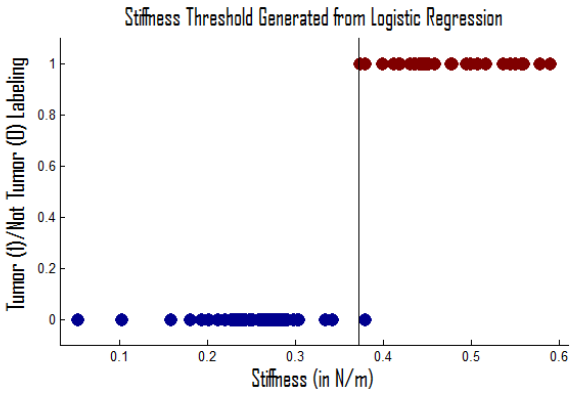

Figure 7: Example of 1-dimensional threshold fitting. The threshold is the black line. For the test sets, all datapoints right of this line would be labeled as tumorous and all points to the left as non-tumorous.

if its stiffness was greater than the threshold, it would be classified as belonging to the tumor. Otherwise, the point would be classified as not belonging to the tumor. See Figure 7 for an example of this 1-Dimensional thresholding.

The threshold was learned using the three different learning algorithms covered in Section 3.2: Logistic Regression, Gaussian Discriminate Analysis, and a Support Vector Machine. Once the palpation data was labeled given the learned stiffness threshold, the supposed center of the tumor was calculated as a mean of all the points classified as a tumor. This concluded the first learning portion of the first method and estimated the center of the tumor; however another learning iteration was performed to determine the size of the tumor.

Towards this end, the palpation data was then concatenated with a new feature: the distance of the palpation point to the supposed center of the tumor. The same three learning algorithms were again utilized to threshold the likely radius of the tumor given our (relabeled) dataset. The data would again be labeled according to this new threshold.

3.4.1 Labeling the Test Set

When implementing this first method on a test set, the stiffness threshold would be used against the test data to provide an initial labeling of the data points. The mean of the data points classified as a tumor served to estimate the center of the tumor. After this initial labeling, LR, GDA, and SVMs were trained to learn the most likely radius of the tumor, and the points were then relabeled according to this threshold. In this method, the stiffness threshold was learned from the training set, while the distance from center of tumor threshold was learned and applied on the test data itself.

3.5 Weighted Clusters followed by Radius of Tumor Learning

The second method first estimated the center of the centroid according to the weighted centroid calculation in Equation 1:

$$center = \frac{1}{n}\sum_{i=0}^{n} 10k_{(i)}^{3}x_{(i)} \qquad \text{(Equation 1)}$$

The parameters {10,3} seen in Equation 1 were tuned by hand and found to give acceptable performance on the training set exhibited by the estimated tumor center being within 2mm of the actual center of the tumor.

Once the center of the centroid was calculated, the palpation data was concatenated with the distance of each palpation from

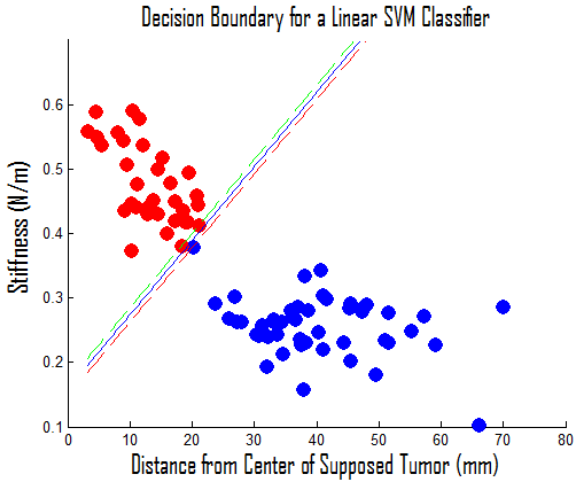

Figure 8: Example of a 2-dimensional threshold fitting. The threshold is the blue line, which the support vector margins are in green and red.

the supposed center of the tumor in the same way as in the method detailed in Section 3.4. This concatenated data and the labeling provided by the training set were used to create a threshold using LR, GDA, and a SVM. This threshold estimates the radius of the tumor.

It is important to know that because this threshold is constructed form the training set, it can only predict tumors of the same size as the training set.

3.5.1 Labeling the Test Set

The estimated center of the tumor was first calculated by Equation 1.Then, each palpation's data was concatenated with its distance from the estimated center. Finally, the points were labeled according to the thresholds developed in training.

3.6 Weighted Clusters followed by Radius of Tumor vs. Stiffness Learning

This method estimated the center of the tumor using the weighted centroid calculation in Equation 1. After the center was calculated and the palpation data was concatenated with the distance from the estimated center of the tumor, a 2-Dimsional plot was created using the data, with one axis being the distance from the estimated center of the tumor, and the other axis being the approximate stiffness observed from the palpation. LR, GDA, and a SVM were again utilized to develop thresholds in this 2-D space. Figure 8 shows the threshold for the training set generated by the SVM.

3.6.1 Labeling the Test Set

The weighted centroid center location was first calculated for the test set, followed by classifying the points based on the thresholds derived in the training phase.

3.7 K-Means Clustering

Shifting away from supervised learning algorithms, this next method of classification first calculated the weighted centroid center location using Equation 1, and generated the same 2-dimensional plot used in Section 3.6. This method would then implement k-means clustering, labeling two clusters to segment the data. The two clusters would classify points as "tumor" and "not tumor".

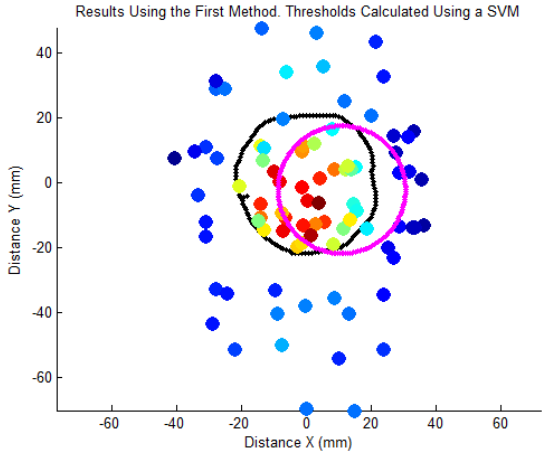

Figure 9: Results from a single trial of the method detailed in Section 3.4 using a SVM for thresholding. The estimated tumor is shown in magenta, the actual tumor is shown in black.

3.8 K-Means Clustering and Radius of Tumor vs. Stiffness Learning

This method I'd like to refer to as "semi-supervised", as it has the notion of a training set but derives thresholds from this "training set" based on an unsupervised labeling of the data. Following the same exact procedure as in 3.7 all the way through labeling the two point classes, the set of data arbitrarily chosen as the "training set" would be used to develop thresholds using LR, GDA, and a SVM in a similar fashion to the method explained in Section 3.6. Recall the "training set" didn't come with a ground-truth labeling, instead the threshold is developed from the K-means clusters labeling of the data.

3.8.1 Labeling the Test Set

Similar to the initialization of labeling the test set in Section 3.5, a 2-D plot was created with distance from the center of the supposed tumor on one axis and stiffness on another. Then, using the thresholds generated form the "semi-training set", the data was classified as part of the tumor or not part of the tumor.

4 Results

4.1 Tumor/Not Tumor Labeling

The three test sets averaged 76 different palpations apiece. The training set has 81 different palpations.

Figure 9 shows the result of one iteration of one learning method graphically. In particular, this analysis was done using the method detailed in section 3.4 and used a SVM to generate all the threshold values.

Figure 10 displays the number of false positives plus false negatives for all methods across all learning algorithms.

4.2 Specific Size and Location Data

The method described in Section 3.4 estimated the center of the supposed tumor to be approximately 1cm away from the actual location.

The methods described in Sections 3.5-3.8 all use a weighted centroid to approximate the center of the supposed tumor using Equation 1. The average error between the actual and estimated center of the tumor was 8.7891mm. All the learning methods appeared to predict the approximate radius of the tumor within 1cm.

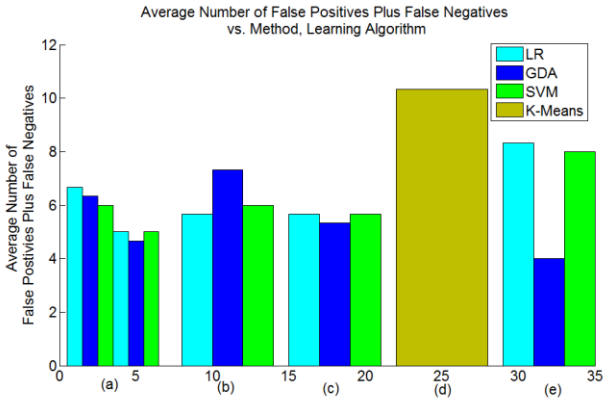

Figure 10: Results indicating the number of false positives plus false negatives against learning methodology, learning algorithm

(a) Stiffness Learning followed by Radius of Tumor Learning (Note since this method labelled the test set points twice, I included errors from both relabelings).
(b) Weighted Clusters followed by Radius of Tumor Learning
(c) Weighted Clusters followed by Radius of Tumor vs. Stiffness Learning
(d) K-Means Clustering
(e) K-Means Clustering and Radius of Tumor vs. Stiffness Learning

5 Discussion

5.1 Number of False Negatives Plus Number of False Positives

Referring to Figure 10, it appears that all the learning methods presented in Section 3 perform reasonably well, although the supervised learning methods did outperform the unsupervised learning methods. GDA tended to be more conservative then LR or a SVM and chose less aggressive thresholds. We can see for 3 of the 4 methods that compared the performance of LR, GDA, and a SVM this conservatism resulted in lower average errors overall (Figure 10a, 10c, 10e).

5.2 Location of Tumor

All the methods suffered a large error with respect to predicting the center of tumor. This is due to two reasons. Firstly, the palpation data wasn't controlled to a specific resolution, so the palpation locations weren't uniformly distributed about the tumor (useful for weighted centroid estimates) and not uniformly distributed within the tumor itself (useful for all methods).

Secondly, the method by which the force data was acquired is too dependent on the configuration of the Phantom robot and isn't uniform over the workspace of the robot. In particular, I was only reading the force data read from the sensor along one axis, regardless of the orientation of the force sensor with respect to the artificial tissue. Towards the edges of the Phantom's workspace, the force sensor is tilted, and the force read by this tilted sensor generated from a palpation normal to the surface of the artificial tissue will not generate forces solely in the z-axis relative to the force-torque sensor.

5.3 Robustness

While all the methods did reasonably well at labeling the points in the testing sets, a major flaw in most of the learning methods was that they are specifically trained to identify a tumor of the same size as the testing set, and are furthermore restricted to identifying only one tumor per test set.

The first method, however, differs from the others in the sense that it thresholds first on solely the stiffness information, regardless of how many tumors are present or the size of tumors. However, the second round of thresholding based on the distance from the estimated center of the tumor assumes one tumor present. Despite this, the increased flexibility of the first method with respect to size of the tumor makes it superior to the other methods explored in this paper in light of this method exhibiting no greater error than the other methods presented.

6 FUTURE WORK

6.1 Hardware

The "artificial tissue" used was a sponge, not a phantom tissue as is desired for palpation experiments. Once a phantom tissue is created, more accurate stiffness data can be attained.

6.2 Data Collection

There are two major areas of improvement for data collection, and one minor. Firstly, with an off-the-shelf path planning algorithm, the Phantom Premium should be able to palpate the artificial tissue autonomously. Secondly, addressing the workspace constraints of the Phantom Premium a more clever way of gathering the force data for palpations taken at the edges of the Phantom Premium's workspace needs to be utilized. This would involve reading the forces on all three axis of the force-torque sensor and given knowledge about the configuration of the sensor relative to the artificial tissue from the Phantom Premium's encoders, transform these forces into the same Cartesian space as the artificial tissue, and record the force vector normal to the surface of the tissue. Additionally, using a non-linear model for tissue stiffness has shown to estimate data taken during palpations of phantom tissues more accurately than the linear model presented [5].

6.3 Learning Algorithms

Ideally, with the autonomous palpation goal mentioned in the hardware section of future work, the learning would happen in real-time such that as the Phantom Premium palpates the artificial tissue sample, it can learn where the tumor is likely located and then gather new data near the thresholds of its 50% confidence region to generate more accurate predictions.

7 CONCLUSIONS

In summary, I've shown that with the data gathered from a force-torque sensor and a robotic device, learning algorithms can be used to estimate the size and location of a tumor-like object embedded into an artificial tissue. Five different learning methods were explored, three with supervised learning and two with unsupervised learning. Although all methods gave comparable results, the first method proposed, one that relabels the test set based on a stiffness threshold learned from the training set, and then with this new labeling develops a threshold for the radius of the is the most robust algorithm, and will be utilized in future works. This work will be extended to include autonomous palpation and real-time learning algorithms.

ACKNOWLEDGEMENTS

I'd like to thank Allison Okamura for her advice throughout this project and trusting me to setup the hardware.

REFERENCES

[1] G. Guthart and J. Salsibury, J.K. The Intuitive TM Telesurgery System: Overview and Application. In IEEE International Conference and Robotics and Automation, Pages 618-621,2000

[2] Gwilliam, James C., Zachary Pezzementi, Erica Jantho, Allison M. Okamua, and Steven Hsiao. "Human vs. Robotic Tactile Sensing: Detecting Lumps in Soft Tissue." *Haptics Symposium* (2010). *IEEE Xplore*. Web. 12 Dec. 2011.

[3] Grant, Michael, and Stephen Boyd. *CVX*. Computer software. *CVX: MATLAB Software for Disciplined Convex Programming*. Vers. 1.21. Web. 11 Dec. 2011. <http://cvxr.com/cvx/>.

[4] S. B. Williams, M. H. Chen, A. V. D'Amico, A. C. Weinberg, R. Kacker, M. S. Hirsch, J. P. Richie, and J. C. Hu. Radical retropubic prostatectomy and robotic-assisted laparoscopic prostatectomy: likelihood of positive surgical margin(s). Urology, 2010 Mar 17. [Epub ahead of print]

[5] Yamamoto, Tomonori. "Applying Tissue Models in Teleoperated Robot-Assisted Surgery" Diss. Johns Hopkins University, Baltimore, 2011. *Dissertations and Theses.* Web.